Logical Methods in Computer Science
Vol. 4 (2:7) 2008, pp. 1–25
www.lmcs-online.org

Submitted Sep. 20, 2007
Published May 15, 2008

ON THE EXPRESSIVENESS AND COMPLEXITY OF ATL

FRANÇOIS LAROUSSINIE [a], NICOLAS MARKEY [b], AND GHASSAN OREIBY [c]

[a] LIAFA, Univ. Paris 7 & CNRS, France
e-mail address: francoisl@liafa.jussieu.fr

[b,c] LSV, ENS Cachan & CNRS, France
e-mail address: {markey,oreiby}@lsv.ens-cachan.fr

Abstract. ATL is a temporal logic geared towards the specification and verification of properties in multi-agents systems. It allows to reason on the existence of strategies for coalitions of agents in order to enforce a given property. In this paper, we first precisely characterize the complexity of ATL model-checking over Alternating Transition Systems and Concurrent Game Structures when the number of agents is not fixed. We prove that it is Δ_2^{P}- and Δ_3^{P}-complete, depending on the underlying multi-agent model (ATS and CGS resp.). We also consider the same problems for some extensions of ATL. We then consider expressiveness issues. We show how ATS and CGS are related and provide translations between these models w.r.t. alternating bisimulation. We also prove that the standard definition of ATL (built on modalities "Next", "Always" and "Until") cannot express the duals of its modalities: it is necessary to explicitely add the modality "Release".

1. Introduction

Model checking. Temporal logics were proposed for the specification of reactive systems almost thirty years ago [CE81, Pnu77, QS82]. They have been widely studied and successfully used in many situations, especially for model checking —the automatic verification that a finite-state model of a system satisfies a temporal logic specification. Two flavors of temporal logics have mainly been studied: *linear-time temporal logics*, *e.g.* LTL [Pnu77], which expresses properties on the possible *executions* of the model; and *branching-time temporal logics*, such as CTL [CE81, QS82], which can express requirements on *states* (which may have several possible futures) of the model.

1998 ACM Subject Classification: F.1.1,F.3.1.
Key words and phrases: multi-agent systems, temporal logic, model checking.
This article is a long version of [LMO07].
[c] This author is sponsored by a PhD grant from Region Île-de-France.

LOGICAL METHODS IN COMPUTER SCIENCE
DOI:10.2168/LMCS-4 (2:7) 2008
© F. Laroussinie, N. Markey, and G. Oreiby
Creative Commons

Alternating-time temporal logic. Over the last ten years, a new flavor of temporal logics has been defined: *alternating-time temporal logics* (ATL) [AHK97]. ATL is a fundamental logic for verifying properties in *synchronous multi-agent systems*, in which several agents can concurrently act upon the behavior of the system. This is particularly interesting for modeling control problems. In that setting, it is not only interesting to know if something *can arrive* or *will arrive*, as can be expressed in CTL or LTL, but rather if some agent(s) can *control* the evolution of the system in order to enforce a given property.

The logic ATL can precisely express this kind of properties, and can for instance state that "there is a strategy for a coalition A of agents in order to eventually reach an accepting state, whatever the other agents do". ATL can be seen as an extension of CTL; its formulae are built on atomic propositions and boolean combinators, and (following the seminal papers [AHK97, AHK98, AHK02]) on modalities $\langle\langle A\rangle\rangle \mathbf{X}\,\varphi$ (coalition A has a strategy to immediately enter a state satisfying φ), $\langle\langle A\rangle\rangle \mathbf{G}\,\varphi$ (coalition A can force the system to always satisfy φ) and $\langle\langle A\rangle\rangle \varphi\,\mathbf{U}\,\psi$ (coalition A has a strategy to enforce $\varphi\,\mathbf{U}\,\psi$).

Multi-agent models. While linear- and branching-time temporal logics are interpreted on Kripke structure, alternating-time temporal logics are interpreted on models that incorporate the notion of *multiple agents*. Two kinds of synchronous multi-agent models have been proposed for ATL in the literature. First *Alternating Transition Systems* (ATSs)[AHK98] have been defined: in any location of an ATS, each agent chooses one *move*, *i.e.*, a subset of locations (the list of possible moves is defined explicitly in the model) in which she would like the execution to go to. When all the agents have made their choice, the intersection of their choices is required to contain one single location, in which the execution enters. In the second family of models, called *Concurrent Game Structures* (CGSs) [AHK02], each of the n agents has a finite number of possible moves (numbered with integers), and, in each location, an n-ary transition function indicates the state to which the execution goes.

Our contributions. First we precisely characterize the complexity of the model checking problem. The original works about ATL provide model-checking algorithms in time $O(m \cdot l)$, where m is the number of transitions in the model, and l is the size of the formula [AHK98, AHK02], thus in PTIME. However, contrary to Kripke structures, the number of transitions in a CGS or in an ATS is not quadratic in the number of states [AHK02], and might even be exponential in the number of agents. PTIME-completeness thus only holds for ATS when the number of agents is bounded, and it is shown in [JD05, JD06] that the problem is strictly[1] harder otherwise, namely NP-hard on ATS and Σ_2^{P}-hard on CGSs where the transition function is encoded as a boolean function. We prove that it is in fact Δ_2^{P}-complete and Δ_3^{P}-complete, resp. We also precisely characterize the complexity of model-checking classical extensions of ATL, depending on the underlying family of models.

Then we address expressiveness questions. First we show how ATSs and CGSs are related by providing translations between these models. Moreover we consider expressiveness questions about ATL modalities. While in LTL and CTL, the dual of "Until" modality can be expressed as a disjunction of "always" and "until", we prove that it is not the case in ATL. In other words, ATL, as defined in [AHK97, AHK98, AHK02], is not as expressive

[1] We adopt the classical hypothesis that the polynomial-time hierarchy does not collapse, and that PTIME $\neq$ NP. We refer to [Pap94] for the definitions about complexity classes, especially about oracle Turing machines and the polynomial-time hierarchy.

as one could expect (while the dual modalities clearly do not increase the complexity of the verification problems).

Related works. In [AHK98, AHK02], ATL has been defined and studied over ATSs and CGSs. In [HRS02], expressiveness issues are considered for ATL* and ATL. Complexity of satisfiability is addressed in [GvD06, WLWW06]. Complexity results about model checking (for ATL, ATL$^+$, ATL*) can be found in [AHK02, Sch04]. Regarding control- and game theory, many papers have focused on this wide area; we refer to [Wal04] for a survey, and to its numerous references for a complete overview.

Plan of the paper. Section 2 contains the formal definitions needed in the sequel. Section 3 deals with the model-checking questions and contains algorithms and complexity analysis for ATSs and CGSs. Section 4 contains our expressiveness results: we first prove that ATSs and CGSs have the same expressive power w.r.t. alternating bisimulation (*i.e.*, any CGS can be translated into an equivalent ATS, and vice-versa). We then present our expressiveness results concerning ATL modalities.

2. Definitions

2.1. **Concurrent Game Structures.** Concurrent game structures are a multi-player extension of classical Kripke structures [AHK02]. Their definition is as follows:

Definition 2.1. A *Concurrent Game Structure* (*CGS* for short) $\mathcal{C}$ is a 6-tuple $(\mathsf{Agt}, \mathsf{Loc}, \mathsf{AP}, \mathsf{Lab}, \mathsf{Mov}, \mathsf{Edg})$ where:

- $\mathsf{Agt} = \{A_1, ..., A_k\}$ is a finite set of *agents* (or *players*);
- Loc and AP are two finite sets of *locations* and *atomic propositions*, resp.;
- $\mathsf{Lab}\colon \mathsf{Loc} \to 2^{\mathsf{AP}}$ is a function labeling each location by the set of atomic propositions that hold for that location;
- $\mathsf{Mov}\colon \mathsf{Loc} \times \mathsf{Agt} \to \mathcal{P}(\mathbb{N}) \setminus \{\varnothing\}$ defines the (finite) set of possible moves of each agent in each location.
- $\mathsf{Edg}\colon \mathsf{Loc} \times \mathbb{N}^k \to \mathsf{Loc}$, where $k = |\mathsf{Agt}|$, is a (partial) function defining the transition table. With each location and each set of moves of the agents, it associates the resulting location.

The intended behaviour is as follows [AHK02]: in a location ℓ, each player A_i chooses one possible move m_{A_i} in $\mathsf{Mov}(\ell, A_i)$ and the next location is given by $\mathsf{Edg}(\ell, m_{A_1}, ..., m_{A_k})$. We write $\mathsf{Next}(\ell)$ for the set of all possible successor locations from ℓ, and $\mathsf{Next}(\ell, A_j, m)$, with $m \in \mathsf{Mov}(\ell, A_j)$, for the restriction of $\mathsf{Next}(\ell)$ to locations reachable from ℓ when player A_j makes the move m.

The way the transition table Edg is encoded has not been made precise in the original definition. Following the remarks of [JD05], we propose two possible encodings:

Definition 2.2.

- An *explicit CGS* is a CGS where the transition table is defined explicitly.

- An *implicit CGS* is a CGS where, in each location ℓ, the transition function is defined by a finite sequence $((\varphi_0, \ell_0), ..., (\varphi_n, \ell_n))$, where $\ell_i \in \mathsf{Loc}$ is a location, and φ_i is a boolean combination of propositions $A_j = c$ that evaluate to true iff agent A_j chooses move c. The transition table is then defined as follows: $\mathsf{Edg}(\ell, m_{A_1}, ..., m_{A_k}) = \ell_j$ iff j is the lowest index s.t. φ_j evaluates to true when players A_1 to A_k choose moves m_{A_1} to m_{A_k}. We require that the last boolean formula φ_n be $\top$, so that no agent can enforce a deadlock.

Besides the theoretical aspect, the implicit description of CGSs may reveal useful in practice, as it allows to not explicitly describe the full transition table.

The size $|\mathcal{C}|$ of a CGS $\mathcal{C}$ is defined as $|\mathsf{Loc}| + |\mathsf{Edg}|$. For explicit CGSs, $|\mathsf{Edg}|$ is the size of the transition table. For implicit CGSs, $|\mathsf{Edg}|$ is the sum of the sizes of the formulas used for the definition of Edg.

2.2. **Alternating Transition Systems.** In the original works about ATL [AHK97], the logic was interpreted on ATSs, which are transition systems slightly different from CGSs:

Definition 2.3. An *Alternating Transition System* (*ATS* for short) $\mathcal{A}$ is a 5-tuple $(\mathsf{Agt}, \mathsf{Loc}, \mathsf{AP}, \mathsf{Lab}, \mathsf{Mov})$ where:

- Agt, Loc, AP and Lab have the same meaning as in CGSs;
- $\mathsf{Mov}\colon \mathsf{Loc} \times \mathsf{Agt} \to \mathcal{P}(\mathcal{P}(\mathsf{Loc}))$ associate with each location ℓ and each agent a the set of possible moves, each move being a subset of Loc. For each location ℓ, it is required that, for any $Q_i \in \mathsf{Mov}(\ell, A_i)$, $\bigcap_{i \leq k} Q_i$ be a singleton.

The intuition is as follows: in a location ℓ, once all the agents have chosen their moves (*i.e.*, a subset of locations), the execution goes to the (only) state that belongs to all the sets chosen by the players. Again $\mathsf{Next}(\ell)$ (resp. $\mathsf{Next}(\ell, A_j, m)$) denotes the set of all possible successor locations (resp. the set of possible successor locations when player A_j chooses the move m).

The size of an ATS is $|\mathsf{Loc}| + |\mathsf{Mov}|$ where $|\mathsf{Mov}|$ is the sum of the number of locations in each possible move of each agent in each location.

We prove in Section 4.1 that CGSs and ATSs have the same expressiveness (w.r.t. alternating bisimilarity [AHKV98]).

2.3. **Coalition, strategy, outcomes of a strategy.** A coalition is a subset of agents. In multi-agent systems, a coalition A plays against its opponent coalition $\mathsf{Agt} \setminus A$ as if they were two single players. We thus extend Mov and Next to coalitions:

- Given $A \subseteq \mathsf{Agt}$ and $\ell \in \mathsf{Loc}$, $\mathsf{Mov}(\ell, A)$ denotes the possible moves for the coalition A from ℓ. Such a move m is composed of a single move for every agent of the coalition, that is $m \stackrel{\text{def}}{=} (m_a)_{a \in A}$. Then, given a move $m' \in \mathsf{Mov}(\ell, \mathsf{Agt} \backslash A)$, we use $m \oplus m'$ to denote the corresponding *complete* move (one for each agent). In ATSs, such a move $m \oplus m'$ corresponds to the unique resulting location; in CGSs, it is given by $\mathsf{Edg}(\ell, m \oplus m')$.
- Next is extended to coalitions in a natural way: given $m = (m_a)_{a \in A} \in \mathsf{Mov}(\ell, A)$, we let $\mathsf{Next}(\ell, A, m)$ denote the restriction of $\mathsf{Next}(\ell)$ to locations reachable from ℓ when every player $A_j \in A$ makes the move m_{A_j}.

Let $\mathcal{S}$ be a CGS or an ATS. A *computation* of $\mathcal{S}$ is an infinite sequence $\rho = \ell_0 \ell_1 \cdots$ of locations such that for any i, $\ell_{i+1} \in \mathsf{Next}(\ell_i)$. We write $\rho[i]$ for the $i+1$-st location ℓ_i. A *strategy* for a player $A_i \in \mathsf{Agt}$ is a function f_{A_i} that maps any finite prefix of a computation

to a possible move for A_i, *i.e.*, satisfying $f_{A_i}(\ell_0 \cdots \ell_m) \in \mathsf{Mov}(\ell_m, A_i)$. A strategy is *state-based* (or *memoryless*) if it only depends on the current state (*i.e.*, $f_{A_i}(\ell_0 \cdots \ell_m) = f_{A_i}(\ell_m)$).

A strategy induces a set of computations from ℓ —called the *outcomes* of f_{A_i} from ℓ and denoted[2] $\mathsf{Out}_{\mathcal{S}}(\ell, f_{A_i})$— that player A_i can enforce: $\ell_0\ell_1 \cdots \in \mathsf{Out}_{\mathcal{S}}(\ell, f_{A_i})$ iff $\ell = \ell_0$ and for any i we have $\ell_{i+1} \in \mathsf{Next}(\ell_i, A_i, f_{A_i}(\ell_0 \cdots \ell_i))$. Given a coalition $A \subseteq \mathsf{Agt}$, a strategy for A is a tuple F_A containing one strategy for each player in A: $F_A = \{f_{A_j} | A_j \in A\}$. The outcomes of F_A from a location ℓ contains the computations enforced by the strategies in F_A: $\ell_0\ell_1 \cdots \in \mathsf{Out}_{\mathcal{S}}(\ell, F_A)$ iff $\ell = \ell_0$ and for any i, $\ell_{i+1} \in \mathsf{Next}(\ell_i, A, (f_a(\ell_0, \cdots, \ell_i))_{a \in A})$. The set of strategies for A is denoted[2] $\mathsf{Strat}_{\mathcal{S}}(A)$. Finally, note that $F_\varnothing$ is empty and $\mathsf{Out}_{\mathcal{S}}(\ell, \varnothing)$ represents the set of all computations from ℓ.

2.4. **The logic ATL.** We now define the logic ATL, whose purpose is to express controllability properties on CGSs and ATSs. Our definition is slightly different from the one proposed in [AHK02]. This difference will be explained and argued in Section 4.2.

Definition 2.4. The syntax of ATL is defined by the following grammar:

$$\begin{array}{rcl} \mathsf{ATL} \ni \varphi_s, \psi_s & ::= & \top \mid P \mid \neg\varphi_s \mid \varphi_s \vee \psi_s \mid \langle\!\langle A \rangle\!\rangle \varphi_p \\ \varphi_p & ::= & \neg\varphi_p \mid \mathbf{X}\,\varphi_s \mid \varphi_s\,\mathbf{U}\,\psi_s \end{array}$$

where P ranges over the set AP and A over the subsets of Agt.

Given a formula $\varphi \in \mathsf{ATL}$, the size of φ, denoted by $|\varphi|$, is the size of the tree representing that formula. The DAG-size of φ is the size of the directed acyclic graph representing that formula (*i.e.*, sharing common subformulas).

In addition, we use standard abbreviations such as $\top$, $\bot$, $\mathbf{F}$, etc. ATL formulae are interpreted over states of a game structure $\mathcal{S}$. The semantics of the main operators is defined as follows[2]:

$$\begin{array}{rcl} \ell \models_{\mathcal{S}} \langle\!\langle A \rangle\!\rangle \varphi_p & \text{iff} & \exists F_A \in \mathsf{Strat}(A).\ \forall \rho \in \mathsf{Out}(\ell, F_A).\ \rho \models_{\mathcal{S}} \varphi_p, \\ \rho \models_{\mathcal{S}} \mathbf{X}\,\varphi_s & \text{iff} & \rho[1] \models_{\mathcal{S}} \varphi_s, \\ \rho \models_{\mathcal{S}} \varphi_s\,\mathbf{U}\,\psi_s & \text{iff} & \exists i.\ \rho[i] \models_{\mathcal{S}} \psi_s \text{ and } \forall 0 \leq j < i.\ \rho[j] \models_{\mathcal{S}} \varphi_s. \end{array}$$

It is well-known that, for the logic ATL, it is sufficient to restrict to state-based strategies (*i.e.*, $\langle\!\langle A \rangle\!\rangle \varphi_p$ is satisfied iff there is a state-based strategy all of whose outcomes satisfy φ_p) [AHK02, Sch04].

Note that $\langle\!\langle \varnothing \rangle\!\rangle \varphi_p$ corresponds to the CTL formula $\mathbf{A}\varphi_p$ (*i.e.*, universal quantification over all computations issued from the current state), while $\langle\!\langle \mathsf{Agt} \rangle\!\rangle \varphi_p$ corresponds to existential quantification $\mathbf{E}\varphi_p$. However, $\neg \langle\!\langle A \rangle\!\rangle \varphi_p$ is generally *not* equivalent to $\langle\!\langle \mathsf{Agt} \setminus A \rangle\!\rangle \neg\varphi_p$ [AHK02, GvD06]: indeed the absence of a strategy for a coalition A to ensure φ does not entail the existence of a strategy for the coalition $\mathsf{Agt}\backslash A$ to ensure $\neg\varphi$. For instance, Fig. 1 displays a (graphical representation of a) 2-player CGS for which, in ℓ_0, both $\neg \langle\!\langle A_1 \rangle\!\rangle \mathbf{X}\, p$ and $\neg \langle\!\langle A_2 \rangle\!\rangle \neg\mathbf{X}\, p$ hold. In such a representation, a transition is labeled with $\langle m_1, m_2 \rangle$ when it corresponds to move m_1 of player A_1 and to move m_2 of player A_2. Fig. 2 represents an "equivalent" ATS with the same property.

[2]We might omit to mention $\mathcal{S}$ when it is clear from the context.

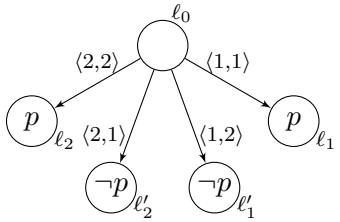

Figure 1: A CGS that is not determined.

$$\mathsf{Loc} = \{\ell_0, \ell_1, \ell_2, \ell'_1, \ell'_2\}$$
$$\mathsf{Mov}(\ell_0, A_1) = \{\{\ell_1, \ell'_1\}, \{\ell_2, \ell'_2\}\}$$
$$\mathsf{Mov}(\ell_0, A_2) = \{\{\ell_1, \ell'_2\}, \{\ell_2, \ell'_1\}\}$$
$$\text{with } \begin{cases} \mathsf{Lab}(\ell_1) = \mathsf{Lab}(\ell_2) = \{p\} \\ \mathsf{Lab}(\ell'_1) = \mathsf{Lab}(\ell'_2) = \varnothing \end{cases}$$

Figure 2: An ATS that is not determined.

3. Complexity of ATL model-checking

In this section, we establish the precise complexity of ATL model-checking. This issue has already been addressed in the seminal papers about ATL, on both ATSs [AHK98] and CGSs [AHK02]. The time complexity is shown to be in $O(m \cdot l)$, where m is the number of transitions and l is the size of the formula. The authors then claim that the model-checking problem is in PTIME (and obviously, PTIME-complete, since it is already for CTL). In fact this only holds for explicit CGSs. In ATSs, the number of transitions might be exponential in the size of the system (more precisely, in the number of agents). This problem —the exponential blow-up of the number of transitions to handle in the verification algorithm— also occurs for implicit CGSs: the standard algorithms running in $O(m \cdot l)$ require exponential time.

Basically, the algorithm for model-checking ATL is similar to that for CTL: it consists in recursively computing fixpoints, based *e.g.* on the following equivalence:

$$\langle\!\langle A \rangle\!\rangle\, p\, \mathbf{U}\, q \equiv \mu Z.(q \vee (p \wedge \langle\!\langle A \rangle\!\rangle\, \mathbf{X}\, Z)) \tag{3.1}$$

The difference with CTL is that we have to deal with the modality $\langle\!\langle A \rangle\!\rangle\, \mathbf{X}$ —corresponding to the *pre-image* of a set of states *for some coalition*— instead of the standard modality EX. In control theory, $\langle\!\langle A \rangle\!\rangle\, \mathbf{X}$ corresponds to the *controllable predecessors* of a set of states for a coalition: $CPre(A, S)$, with $A \subseteq \mathsf{Agt}$ and $S \subseteq \mathsf{Loc}$, is defined as follows:

$$CPre(A, S) \stackrel{\text{def}}{=} \{\ell \in \mathsf{Loc} \mid \exists m_A \in \mathsf{Mov}(\ell, A) \text{ s.t. } \mathsf{Next}(\ell, A, m_A) \subseteq S\}$$

The crucial point of the model-checking algorithm is the computation of the set $CPre(A, S)$.

In the sequel, we establish the exact complexity of computing *CPre* (more precisely, given $A \subset \mathsf{Agt}$, $S \subseteq \mathsf{Loc}$, and $\ell \in \mathsf{Loc}$, the complexity of deciding whether $\ell \in CPre(A, S)$), and of ATL model-checking for our three kinds of multi-agent systems.

3.1. Model checking ATL on explicit CGSs.

As already mentionned, the precise complexity of ATL model-checking over explicit CGSs was established in [AHK02]:

Theorem 3.1. *ATL model-checking over explicit CGSs is* PTIME*-complete.*

To our knowledge, the precise complexity of computing *CPre* in explicit CGSs has never been considered. The best upper bound is PTIME, which is sufficient for deriving the PTIME complexity of ATL model-checking.

In fact, given a location ℓ, a set of locations S and a coalition A, deciding whether $\ell \in CPre(A, S)$ has complexity much lower than PTIME:

Proposition 3.2. *Computing CPre in explicit CGSs is in* AC^0.

Proof. We begin with precisely defining how the input is encoded as a sequence of bits:

- the first $|\mathsf{Agt}|$ bits define the coalition: the i-th bit is a 1 iff agent A_i belongs to A;
- the following $|\mathsf{Loc}|$ bits of input define the set S;
- for the sake of simplicity, we assume that all the agents have the same number of moves in ℓ. We write p for that number, which we assume is at least 2. The transition table $\mathsf{Edg}(\ell)$ is then given as a sequence of p^k sets of $\log(|\mathsf{Loc}|)$ bits.

As a first step, it is rather easy to modify the input in order to have the following form:

- first the k bits defining the coalition;
- then, a sequence of p^k bits defining whether the resulting state belongs to S.

This is achieved by p^k copies of the same AC^0 circuit.

We now have to build a circuit that will "compute" whether coalition A has a strategy for ending up in S. Since circuits must only depend on the size of the input, we cannot design a circuit for coalition A. Instead, we build one circuit for each possible coalition (their number is exponential in the number of agents, but polynomial in the size of the input, provided that $p \geq 2$), and then select the result corresponding to coalition A.

Thus, for each possible coalition B, we build one circuit whose final node will evaluate to 1 iff $\ell \in \mathit{CPre}(B, S)$. This is achieved by an unbounded fan-in circuit of depth 2: at the first level, we put $p^{|B|}$ AND-nodes, representing each of the $p^{|B|}$ possible moves for coalition B. Each of those nodes is linked to $p^{k-|B|}$ bits of the transition table, corresponding to the set of possible $p^{k-|B|}$ moves of the opponents. At the second level, an OR-node is linked to all the nodes at depth 1.

Clearly enough, the OR-node at depth 2 evaluates to true iff coalition B has a strategy to reach S. Moreover, there are $\binom{k}{l}$ coalitions of size l, each of which is handled by a circuit of $p^l + 1$ nodes. The resulting circuit thus has $(p+1)^k + 2^k$ nodes, which is polynomial in the size of the input. This circuit is thus an AC^0 circuit.

It simply remains to return the result corresponding to the coalition A. This is easily achieved in AC^0. □

3.2. **Model checking ATL on implicit CGSs.** Assuming that the transitions issued from ℓ are given —in the transition table— by the sequence $((\varphi_0, \ell_0), (\varphi_1, \ell_1), \ldots, (\varphi_n, \ell_n))$, we have: $\ell \in \mathit{CPre}(A, S)$ iff there exists $m_A \in \mathsf{Mov}(\ell, A)$, s.t. there is no $m_{\bar{A}} \in \mathsf{Mov}(\ell, \mathsf{Agt}\backslash A)$ and $\ell_i \in \mathsf{Loc}\backslash S$ s.t.[3] $\varphi_i[m_A \oplus m_{\bar{A}}] \equiv \top$ and $\varphi_j[m_A \oplus m_{\bar{A}}] \equiv \bot$ for any $j < i$. Thus we look for a move $m_A \in \mathsf{Mov}(\ell, A)$ s.t. for all $m_{\bar{A}} \in \mathsf{Mov}(\ell, \bar{A})$, the negation of $\bigvee_{\ell_i \in \mathsf{Loc}\backslash S}(\varphi_i[m_A] \wedge \bigwedge_{j<i} \neg\varphi_j[m_A])$ holds.

This problem corresponds to an instance of the Σ_2^{P}-complete problem EQSAT_2:

EQSAT_2:

> **Input::** two families of variables $X = \{x^1, ..., x^n\}$ and $Y = \{y^1, ..., y^n\}$, a boolean formula φ on the set of variables $X \cup Y$.
> **Output::** True iff $\exists X.\ \forall Y.\ \varphi$.

And indeed, as a direct corollary of [JD05, Lemma 1], we have:

[3] Given $m = (m_a)_{a \in A}$ for $A \subseteq \mathsf{Agt}$, $\varphi[m]$ denotes the formula where every proposition "$A_j = c$" with $A_j \in A$ is replaced by $\top$ if $m_{A_j} = c$, and by $\bot$ otherwise. If $A = \mathsf{Agt}$, $\varphi[m]$ is boolean expression.

Proposition 3.3. *Computing CPre in implicit CGSs is* Σ_2^{P}*-complete.*

Proof. The membership in Σ_2^{P} follows directly the above remarks. A Σ_2^{P} procedure is explicitly described in Algorithm 1.

Procedure `co-strategy`$(q, (\varphi_i, \ell_i)_i, (m_a)_{a\in A}, S)$
`//checks if the opponents have a co-strategy to` $(m_a)_{a\in A}$ `to avoid` S
begin
  foreach $\bar{a} \in \bar{A}$ **do**
    $m_{\bar{a}} \leftarrow guess(q, \bar{a})$;
  $i \leftarrow 0$;
  while $\neg\varphi_i(m_a, m_{\bar{a}})$ **do**
    $i \leftarrow i + 1$;
  if $\ell_i \notin S$ **then**
    return *yes*;
  else
    return *no*;
end
Procedure `CPre`(A, S) **begin**
  $W \leftarrow \varnothing$;
  foreach $q \in \mathcal{C}$ **do**
    foreach $a \in A$ **do**
      $m_a \leftarrow guess(q, a)$;
    if *not* `co-strategy`$(q, (\varphi_i, \ell_i)_i, (m_a)_{a\in A}, S)$ **then**
      $W \leftarrow W \cup \{q\}$;
  return W;
end

Algorithm 1: Computing *CPre* on implicit CGS.

Concerning hardness in Σ_2^{P}, we directly use the construction of [JD05, Lemma 1]: from an instance $\exists X.\ \forall Y.\ \varphi$ of EQSAT_2, one consider an implicit CGS with three states q_1, $q_\top$ and $q_\bot$, and $2n$ agents A^1, ..., A^n, B^1, ..., B^n, each having two possible choices in q_1 and only one choice in $q_\top$ and $q_\bot$. The transitions out of $q_\top$ and $q_\bot$ are self-loops. The transitions from q_1 are given by: $\delta(q_1) = ((\varphi[x^j \leftarrow (A^j \stackrel{?}{=} 1), y^j \leftarrow (B^j \stackrel{?}{=} 1)], q_\top)(\top, q_\bot))$.

Then clearly, q_1 belongs to $CPre(\{A^1, ..., A^n\}, \{q_\top\})$ iff there exists a valuation for variables in X s.t. φ is true whatever B-agents choose for Y. □

The complexity of ATL model checking over implicit CGS is higher: the proof of Σ_2^{P}-hardness of $CPre(A, S)$ can easily be adapted to prove Π_2^{P}-hardness. Indeed consider the dual (thus Π_2^{P}-complete) problem AQSAT_2, in which, with the same input, the output is the value of $\forall X.\ \exists Y.\ \varphi$. Then it suffices to consider the same implicit CGS, and the formula $\neg \langle\!\langle A^1, ..., A^n \rangle\!\rangle \mathbf{X} \neg q_\top$. It states that there is no strategy for players A^1 to A^n to avoid $q_\top$: whatever their choice, players B^1 to B^n can enforce φ.

This contradicts the claim in [JD05] that model checking ATL would be Σ_2^{P}-complete. In fact there is a flaw in their algorithm about the way it handles negation (and indeed their result holds only for the *positive* fragment of ATL [JD08]): games played on CGSs (and ATSs) are generally not determined, and the fact that a player has no strategy to

enforce φ does not imply that the other players have a strategy to enforce $\neg\varphi$. It rather means that the other players have a *co-strategy* to enforce $\neg\varphi$ (by a *co-strategy*, we mean a way to react to each move of their opponents [GvD06]).

Still, using the expression of ATL modalities as fixpoint formulas (see Eq. (3.1)), we can compute the set of states satisfying an ATL formula by a polynomial number of computations of *CPre*, which yields a Δ_3^{P} algorithm:

Proposition 3.4. *Model checking ATL on implicit CGSs is in* Δ_3^{P}.

Note that, since the algorithm consists in labeling the locations with the subformulae it satisfies, that complexity holds even if we consider the DAG-size of the formula.

To prove hardness in Δ_3^{P}, we introduce the following Δ_3^{P}-complete problem [LMS01, Sch04].

SNSAT$_2$:

Input:: m families of variables $X_i = \{x_i^1, ..., x_i^n\}$, m families of variables $Y_i = \{y_i^1, ..., y_i^n\}$, m variables z_i, m boolean formulae φ_i, with φ_i involving variables in $X_i \cup Y_i \cup \{z_1, ..., z_{i-1}\}$.

Output:: The value of z_m, defined by

$$\begin{cases} z_1 \stackrel{\text{def}}{=} \exists X_1.\ \forall Y_1.\ \varphi_1(X_1, Y_1) \\ z_2 \stackrel{\text{def}}{=} \exists X_2.\ \forall Y_2.\ \varphi_2(z_1, X_2, Y_2) \\ \quad\ldots \\ z_m \stackrel{\text{def}}{=} \exists X_m.\ \forall Y_m.\ \varphi_m(z_1, ..., z_{m-1}, X_m, Y_m) \end{cases}$$

And we have:

Proposition 3.5. *Model checking ATL on implicit CGSs is* Δ_3^{P}*-hard.*

Proof. We pick an instance $\mathcal{I}$ of SNSAT$_2$, and reduce it to an instance of the ATL model-checking problem. Note that such an instance uniquely defines the values of variables z_i. We write $v_{\mathcal{I}} \colon \{z_1, ..., z_m\} \to \{\top, \bot\}$ for this valuation. Also, when $v_{\mathcal{I}}(z_i) = \top$, there exists a witnessing valuation for variables in X_i. We extend $v_{\mathcal{I}}$ to $\{z_1, ..., z_m\} \cup \bigcup_i X_i$, with $v_{\mathcal{I}}(x_i^j)$ being a witnessing valuation if $v_{\mathcal{I}}(z_i) = \top$.

We now define an implicit CGS $\mathcal{C}$ as follows: it contains mn agents A_i^j (one for each x_i^j), mn agents B_i^j (one for each y_i^j), m agents C_i (one for each z_i), and one extra agent D. The structure is made of m states q_i, m states $\overline{q_i}$, m states s_i, and two states $q_\top$ and $q_\bot$. There are three atomic propositions: $s_\top$ and $s_\bot$, that label the states $q_\top$ and $q_\bot$ resp., and an atomic proposition s labeling states s_i. The other states carry no label.

Except for D, the agents represent booleans, and thus always have two possible choices (0 and 1). Agent D always has m choices (0 to $m-1$). The transition relation is defined as follows: for each i,

$$\begin{array}{l} \delta(\overline{q_i}) = ((\top, s_i)); \\ \delta(s_i) = ((\top, q_i)); \\ \delta(q_\top) = ((\top, q_\top)); \\ \delta(q_\bot) = ((\top, q_\bot)); \end{array} \qquad \delta(q_i) = \left(\begin{array}{l} ((D \stackrel{?}{=} 0) \wedge \varphi_i[x_i^j \leftarrow (A_i^j \stackrel{?}{=} 1), \\ \qquad y_i^j \leftarrow (B_i^j \stackrel{?}{=} 1), z_k \leftarrow (C_k \stackrel{?}{=} 1)], q_\top) \\ ((D \stackrel{?}{=} 0), q_\bot) \\ ((D \stackrel{?}{=} k) \wedge (C_k \stackrel{?}{=} 1), q_k) \text{ for each } k < i \\ ((D \stackrel{?}{=} k) \wedge (C_k \stackrel{?}{=} 0), \overline{q_k}) \text{ for each } k < i \\ (\top, q_\top) \end{array} \right)$$

Intuitively, from state q_i, the boolean agents chose a valuation for the variable they represent, and agent D can either choose to check if the valuation really witnesses φ_i (by choosing move 0), or "challenge" player C_k, with move $k < i$.

The ATL formula is built recursively as follows:

$$\psi_0 \stackrel{\text{def}}{=} \top$$

$$\psi_{k+1} \stackrel{\text{def}}{=} \langle\langle \mathsf{AC} \rangle\rangle (\neg s)\, \mathbf{U}\, (q_\top \vee \mathbf{EX}\, (s \wedge \mathbf{EX}\, \neg\psi_k))$$

where AC stands for the coalition $\{A_1^1, ..., A_m^n, C_1, ..., C_m\}$.

Let $f_{\mathcal{I}}(A)$ be the state-based strategy for agent $A \in \mathsf{AC}$ that consists in playing according to the valuation $v_{\mathcal{I}}$ (*i.e.* move 0 if the variable associated with A evaluates to 0 in $v_{\mathcal{I}}$, and move 1 otherwise). The following lemma completes the proof of Proposition 3.5:

Lemma 3.6. *For any $i \leq m$ and $k \geq i$, the following three statements are equivalent:*
(a) $\mathcal{C}, q_i \models \psi_k$*;*
(b) *the strategies $f_{\mathcal{I}}$ witness the fact that $\mathcal{C}, q_i \models \psi_k$;*
(c) *variable z_i evaluates to $\top$ in $v_{\mathcal{I}}$.*

Proof. Clearly, *(b)* implies *(a)*. We prove that *(a)* implies *(c)* and that *(c)* implies *(b)* by induction on i.

First assume that $q_1 \models \psi_j$, for some $j \geq 1$. Since only $q_\top$ and $q_\bot$ are reachable from q_1, we have $q_1 \models \langle\langle \mathsf{AC} \rangle\rangle \mathbf{X}\, q_\top$. We are (almost) in the same case as in the Σ_2^{P} reduction of [JD05]: there is a valuation of the variables x_1^1 to x_1^n s.t., whatever players D and B_1^1 to B_m^n decide, the run will end up in $q_\top$. This holds in particular if player D chooses move 0: for any valuation of the variables y_1^1 to y_1^n, $\psi_1(X_1, Y_1)$ holds true, and z_1 evaluates to true in $v_{\mathcal{I}}$.

Secondly, if z_1 evaluates to true, then $v_{\mathcal{I}}(x_1^1)$, ..., $v_{\mathcal{I}}(x_1^n)$ are such that, whatever the value of y_1^1 to y_1^n, ψ_1 holds true. If players A_1^1 to A_1^n play according to $f_{\mathcal{I}}$, then players D and B_1^1 to B_1^n cannot avoid state $q_\top$, and $q_1 \models \langle\langle \mathsf{AC} \rangle\rangle \mathbf{X}\, q_\top$, thus also ψ_k when $k \geq 1$.

We now assume the result holds up to index $i \geq 1$, and prove that it also holds at step $i+1$. Assume $q_{i+1} \models \psi_{k+1}$, with $k \geq i$. There exists a strategy witnessing ψ_{k+1}, *i.e.*, s.t. all the outcomes following this strategy satisfy $(\neg s)\, \mathbf{U}\, (q_\top \vee \mathbf{EX}\, (s \wedge \mathbf{EX}\, \neg\psi_k))$. Depending on the move of player D in state q_{i+1}, we get several informations: first, if player D plays move l, with $1 \leq l \leq i$, the play goes to state q_l or $\overline{q_l}$, depending on the choice of player C_l.

- if player C_l chose move 0, the run ends up in $\overline{q_l}$. Since the only way out of that state is to enter state s_l, labeled by s, we get that $\overline{q_l} \models \mathbf{EX}\, (s \wedge \mathbf{EX}\, \neg\psi_k)$, *i.e.*, that $q_l \models \neg\psi_k$. By i.h., we get that z_l evaluates to false in our instance of SNSAT_2.
- if player C_l chose move 1, the run goes to q_l. In that state, players in AC can keep on applying their strategy, which ensures that $q_l \models \psi_{k+1}$, and, by i.h., that z_l evaluates to true in $\mathcal{I}$.

Thus, the strategy for AC to enforce ψ_{k+1} in q_{i+1} requires players C_1 to C_i to play according to $v_{\mathcal{I}}$ and the validity of these choices can be verified by the "opponent" D.

Now, if player D chooses move 0, all the possible outcomes will necessarily immediately go to $q_\top$ (since ψ_{k+1} holds, and since $q_\bot \not\models \mathbf{EX}\, (s \wedge \mathbf{EX}\, \neg\psi_k)$). We immediately get that players B_{i+1}^1 to B_{i+1}^n cannot make ψ_{i+1} false, hence that z_{i+1} evaluates to true in $\mathcal{I}$.

Secondly, if z_{i+1} evaluates to true, assume players in AC play according to $f_{\mathcal{I}}$, and consider the possible moves of player D:

- if player D chooses move 0, since z_{i+1} evaluates to true and since players C_1 to C_i and A^1_{i+1} to A^n_{i+1} have played according $v_{\mathcal{I}}$, there is no way for player B^1_{i+1} to B^n_{i+1} to avoid state $q_\top$.
- if player D chooses some move l between 1 and i, the execution will go into state q_l or $\overline{q_l}$, depending on the move of C_l.
  - if C_l played move 0, *i.e.*, if z_l evaluates to false in $v_{\mathcal{I}}$, the execution goes to state $\overline{q_l}$, and we know by i.h. that $q_l \models \neg\psi_k$. Thus $\overline{q_l} \models \mathbf{EX}\,(s \wedge \mathbf{EX}\,\neg\psi_k)$, and the strategy still fulfills the requirement.
  - if C_l played move 1, *i.e.*, if z_l evaluates to true, then the execution ends up in state q_l, in which, by i.h., the strategy $f_{\mathcal{I}}$ enforces ψ_{k+1}.
- if player D plays some move l with $l > i$, the execution goes directly to $q_\top$, and the formula is fulfilled. □

With Proposition 3.4, this implies:

Theorem 3.7. *Model checking ATL on implicit CGSs is Δ_3^{P}-complete.*

3.3. **Model checking ATL on ATSs.** For ATSs also, computing *CPre* (and thus model-checking ATL) cannot be achieved in PTIME. A direct corollary of [JD05, Lemma 4] is:

Proposition 3.8. *Computing CPre in ATSs is* NP*-complete.*

Proof. Algorithm 2 shows how to compute *CPre* in NP in ATSs: it amounts to guessing a move for each player in the coalition, and to check whether the resulting possible next states are all in S.

```
Procedure CPre(A, S) begin
    W ← ∅;
    foreach q ∈ C do
        foreach a ∈ A do
            //Guess a move for player a from a state q
            m_a ← guess(q, a);
        if ⋂_{a∈A} m_a ⊆ S then
            W ← W ∪ {q};
    return W;
end
```

Algorithm 2: Computing *CPre* for ATS

Again, NP-hardness follows from [JD05, Lemma 4]. We propose here a slightly different proof, that will be a first step to the Δ_2^{P}-hardness proof below.

The proof is a direct reduction from 3SAT: let $\mathcal{I} = (S^1, ..., S^n)$ be an instance of 3SAT over variables $X = \{x^1, ..., x^m\}$. We assume that $S^j = \alpha^{j,1}s^{j,1} \vee \alpha^{j,2}s^{j,2} \vee \alpha^{j,3}s^{j,3}$ where $s^{j,k} \in X$ and $\alpha^{j,k} \in \{0, 1\}$ indicates whether variable $s^{j,k}$ is taken negatively (0) or positively (1). We assume without loss of generality that no clauses contain both one proposition and its negation.

With such an instance, we associate the following ATS $\mathcal{A}$. It contains $8n+1$ states: one state q, and, for each clause S^j, eight states $q^{j,0}$ to $q^{j,7}$. Intuitively, the state $q^{j,k}$ corresponds

to a clause $B^{j,k} = k_1 s^{j,1} \vee k_2 s^{j,2} \vee k_3 s^{j,3}$, where $k_1 k_2 k_3$ corresponds to the binary notation for k. There is only one atomic proposition α in our ATS: a state $q^{j,k}$ is labeled with α iff it does not correspond to clause S^j. By construction, for each j, only one of the states $q^{j,0}$ to $q^{j,7}$ is not labeled with α.

There are $m+1$ players, where m is the number of variables that appear in $\mathcal{I}$. With each x^i is associated a player A^i. The extra player is named D. Only the transitions from q are relevant for this reduction. We may assume that the other states only carry a self-loop. In q, player A^i decides the value of x^i. She can thus choose between two sets of next states, namely the states corresponding to clauses that are not made true by her choice:

$$\{q^{j,k} \mid \forall l \leq 3.\ s^{j,l} \neq x^i \text{ or } \alpha^{i,l} = 0\} \qquad \text{if } x^i = \top$$
$$\{q^{j,k} \mid \forall l \leq 3.\ s^{j,l} \neq x^i \text{ or } \alpha^{i,l} = 1\} \qquad \text{if } x^i = \bot$$

Last, player D has n choices, namely $\{q^{1,0}, ..., q^{1,7}\}$ to $\{q^{n,0}, ..., q^{n,7}\}$.

We first prove the singleton requirement for ATSs' transitions: the intersections of the choices of the agents must be a singleton. Once players A^1 to A^m have chosen their moves, all the variables have been assigned a value. Under that valuation, for each $j \leq n$, exactly one clause among $B^{j,0}$ to $B^{j,7}$ evaluates to false (thanks to our requirement that a literal cannot appear together with its negation in the same clause). Intersecting with the choice of player D, we end up with one single state (corresponding to the only clause, among those chosen by D, that evaluates to false).

Now, let $\varphi = \langle\!\langle A^1, ..., A^m \rangle\!\rangle \mathbf{X}\, \alpha$. That $q \models \varphi$ indicates that players A^1 to A^m can choose a valuation for x^1 to x^m s.t. player D will not be able to find a clause of the original instance (*i.e.*, not labeled with α) that evaluates to false (*i.e.*, that is not made true by any of the choices of the players A^1 to A^m). In that case, the instance is satisfiable. Conversely, if the instance is satisfiable, it suffices for the players A^1 to A^m to play according to a satisfying valuation of variables x^1 to x^m. Since this valuation makes all the original clauses true, it yields a strategy that only leads to states labeled with α. □

As in the case of implicit CGSs, we combine the fixpoint expressions of ATL modalities together with the NP algorithm for computing *CPre*. This yield a Δ_2^{P} algorithm for full ATL:

Proposition 3.9. *Model checking ATL over ATSs is in* Δ_2^{P}.

This turns out to be optimal:

Proposition 3.10. *Model checking ATL on ATSs is* Δ_2^{P}*-hard.*

Proof. The proof is by a reduction of the Δ_2^{P}-complete problem SNSAT [LMS01]:

SNSAT:

Input:: p families of variables $X_r = \{x_r^1, ..., x_r^m\}$, p variables z_r, p boolean formulae φ_r in 3-CNF, with φ_r involving variables in $X_r \cup \{z_1, ..., z_{r-1}\}$.

Output:: The value of z_p, defined by

$$\begin{cases} z_1 \stackrel{\text{def}}{=} \exists X_1.\ \varphi_1(X_1) \\ z_2 \stackrel{\text{def}}{=} \exists X_2.\ \varphi_2(z_1, X_2) \\ z_3 \stackrel{\text{def}}{=} \exists X_3.\ \varphi_3(z_1, , z_2, X_3) \\ \quad\quad \cdots \\ z_p \stackrel{\text{def}}{=} \exists X_p.\ \varphi_p(z_1, ..., z_{p-1}, X_p) \end{cases}$$

Let $\mathcal{I}$ be an instance of SNSAT, where we assume that each φ_r is made of n clauses S_r^1 to S_r^n, with $S_r^j = \alpha_r^{j,1} s_r^{j,1} \vee \alpha_r^{j,2} s_r^{j,2} \vee \alpha_r^{j,3} s_r^{j,3}$. Again, such an instance uniquely defines a valuation $v_{\mathcal{I}}$ for variables z_1 to z_r, that can be extended to the whole set of variables by choosing a witnessing valuation for x_r^1 to x_r^n when z_r evaluates to true.

We now describe the ATS $\mathcal{A}$: it contains $(8n+3)p$ states:

- p states $\overline{q_r}$ and p states q_r,
- p states s_r,
- and for each formula φ_r, for each clause S_r^j of φ_r, eight states $q_r^{j,0}, ..., q_r^{j,7}$, as in the previous reduction.

States s_r are labelled with the atomic proposition s, and states $q_r^{j,k}$ that do not correspond to clause S_r^j are labeled with α.

There is one player A_r^j for each variable x_r^j, one player C_r for each z_r, plus one extra player D. As regards transitions, there are self-loops on each state $q_r^{j,k}$, single transitions from each $\overline{q_r}$ to the corresponding s_r, and from each s_r to the corresponding q_r. From state q_r,

- player A_r^j will choose the value of variable x_r^j, by selecting one of the following two sets of states:

$$\{q_r^{g,k} \mid \forall l \leq 3.\ s_r^{g,l} \neq x_r^j \text{ or } \alpha_r^{g,l} = 0\} \cup \{q_t, \overline{q_t} \mid t < r\} \qquad \text{if } x_r^j = \top$$
$$\{q_r^{g,k} \mid \forall l \leq 3.\ s_r^{g,l} \neq x_r^j \text{ or } \alpha_r^{g,l} = 1\} \cup \{q_t, \overline{q_t} \mid t < r\} \qquad \text{if } x_r^j = \bot$$

  Both choices also allow to go to one of the states q_t or $\overline{q_t}$. In q_r, players A_t^j with $t \neq r$ have one single choice, which is the whole set of states.
- player C_t also chooses for the value of the variable it represents. As for players A_r^j, this choice will be expressed by choosing between two sets of states corresponding to clauses that are not made true. But as in the proof of Prop. 3.5, players C_t will also offer the possibility to "verify" their choice, by going either to state q_t or $\overline{q_t}$. Formally, this yields two sets of states:

$$\{q_r^{g,k} \mid \forall l \leq 3.\ s_r^{g,l} \neq z_t \text{ or } \alpha_r^{g,l} = 0\} \cup \{q_u, \overline{q_u} \mid u \neq t\} \cup \{q_t\} \quad \text{if } z_t = \top$$
$$\{q_r^{g,k} \mid \forall l \leq 3.\ s_r^{g,l} \neq z_t \text{ or } \alpha_r^{g,l} = 1\} \cup \{q_u, \overline{q_u} \mid u \neq t\} \cup \{\overline{q_t}\} \quad \text{if } z_t = \bot$$

- Last, player D chooses either to challenge a player C_t, with $t < r$, by choosing the set $\{q_t, \overline{q_t}\}$, or to check that a clause S_r^j is fulfilled, by choosing $\{q_r^{j,0}, ..., q_r^{j,7}\}$.

Let us first prove that any choices of all the players yields exactly one state. It is obvious except for states q_r. For a state q_r, let us first restrict to the choices of all the players A_r^j and C_r, then:

- if we only consider states $q_r^{1,0}$ to $q_r^{n,7}$, the same argument as in the previous proof ensures that precisely on state per clause is chosen,
- if we consider states q_t and $\overline{q_t}$, the choices of players B_t ensure that exactly one state has been chosen in each pair $\{q_t, \overline{q_t}\}$, for each $t < r$.

Clearly, the choice of player D will select exactly one of the remaining states.

Now, we build the ATL formula. It is a recursive formula (very similar to the one used in the proof of Prop. 3.5), defined by $\psi_0 = \top$ and (again writing AC for the set of players $\{A_1^1, ..., A_p^m, C_1, ..., C_p\}$):

$$\psi_{r+1} \stackrel{\text{def}}{=} \langle\langle \mathsf{AC} \rangle\rangle (\neg s)\, \mathbf{U}\, (\alpha \vee \mathbf{EX}\, (s \wedge \mathbf{EX}\, \neg \psi_r)).$$

Then, writing $f_{\mathcal{I}}$ for the state-based strategy associated to $v_{\mathcal{I}}$:

Lemma 3.11. *For any $r \leq p$ and $t \geq r$, the following statements are equivalent:*
(a) $q_r \models \psi_t$;
(b) the strategies $f_{\mathcal{I}}$ witness the fact that $q_r \models \psi_t$;
(c) variable z_r evaluates to true in $v_{\mathcal{I}}$.

Proof. We prove by induction on r that *(a)* implies *(c)* and that *(c)* implies *(b)*, the last implication being obvious. For $r = 1$, since no s-state is reachable, it amounts to the previous proof of NP-hardness.

Assume the result holds up to index r. Then, if $q_{r+1} \models \psi_{t+1}$ for some $t \geq r$, we pick a strategy for coalition AC witnessing this property. Again, we consider the different possible choices available to player D:

- if player D chooses to go to one of q_u and $\overline{q_u}$, with $u < r+1$: the execution ends up in q_u if player C_u chose to set z_u to true. But in that case, formula ψ_{t+1} still holds in q_u, which yields by i.h. that z_u really evaluates to true in $v_{\mathcal{I}}$. Conversely, the execution ends up in $\overline{q_u}$ if player C_u set z_u to false. In that case, we get that $q_u \models \neg\psi_t$, with $t \geq u$, which entails by i.h. that z_u evaluates to false.

  This first case entails that player C_1 to C_r chose the correct value for variables z_1 to z_r.
- if player D chooses a set of eight states corresponding to a clause S^j_{r+1}, then the strategy of other players ensures that the execution will reach a state labeled with α. As in the previous reduction, this indicates that the corresponding clause has been made true by the choices of the other players.

Putting all together, this proves that variable z_{r+1} evaluates to true.

Now, if variable z_{r+1} evaluates to true, Assume the players in AC play according to valuation $f_{\mathcal{I}}$. Then

- if player D chooses to go to a set of states that correspond to a clause of φ_{r+1}, he will necessarily end up in a state that is labeled with α, since the clause is made true by the valuation we selected.
- if player D chooses to go to one of q_u or $\overline{q_u}$, for some u, then he will challenge player B_u to prove that his choice was correct. By i.h., and since player B_u played according to $f_{\mathcal{I}}$, formula $(\neg s)\,\mathbf{U}\,(\alpha \vee \mathbf{EX}\,(s \wedge \mathbf{EX}\,\neg\psi_{t+1}))$ will be satisfied, for any $t \geq u$. □

We end up with the precise complexity of ATL model-checking on ATSs:

Theorem 3.12. *Model checking ATL on ATSs is Δ_2^{P}-complete.*

3.4. **Beyond ATL.** As for classical branching-time temporal logics, we can consider several extensions of ATL by allowing more possibilities in the way of combining quantifiers over strategies and temporal modalities. We define ATL* [AHK02] as follows:

Definition 3.13. The syntax of ATL* is defined by the following grammar:

$$\begin{array}{rcl} \mathsf{ATL}^\star \ni \varphi_s, \psi_s & ::= & \top \mid P \mid \neg\varphi_s \mid \varphi_s \vee \psi_s \mid \langle\!\langle A \rangle\!\rangle\, \varphi_p \\ \varphi_p, \psi_p & ::= & \varphi_s \mid \neg\varphi_p \mid \varphi_p \vee \psi_p \mid \mathbf{X}\,\varphi_p \mid \varphi_p\,\mathbf{U}\,\psi_p \end{array}$$

where P and A range over AP and 2^{Agt}, resp.

The size and DAG-size of an $\mathsf{ATL}^\star$ formula are defined in the same way as for ATL. $\mathsf{ATL}^\star$ formulae are interpreted over states of a game structure $\mathcal{S}$, the semantics of the main modalities is as follows (if $\rho = \ell_0\,\ell_1\,\ldots$, we write ρ^i for the $i+1$-st suffix, starting from ℓ_i):

$$\begin{aligned}
\ell \models_{\mathcal{S}} \langle\!\langle A\rangle\!\rangle\, \varphi_p \quad &\text{iff} \quad \exists F_A \in \mathsf{Strat}(A).\ \forall \rho \in \mathsf{Out}(\ell, F_A).\ \rho \models_{\mathcal{S}} \varphi_p,\\
\rho \models_{\mathcal{S}} \varphi_s \quad &\text{iff} \quad \rho[0] \models_{\mathcal{S}} \varphi_s\\
\rho \models_{\mathcal{S}} \mathbf{X}\,\varphi_p \quad &\text{iff} \quad \rho^1 \models_{\mathcal{S}} \varphi_p,\\
\rho \models_{\mathcal{S}} \varphi_p\,\mathbf{U}\,\psi_p \quad &\text{iff} \quad \rho^i \models_{\mathcal{S}} \psi_p \text{ and } \forall 0 \leq j < i.\ \rho^j \models_{\mathcal{S}} \varphi_p
\end{aligned}$$

ATL is the fragment of $\mathsf{ATL}^\star$ where each modality $\mathbf{U}$ or $\mathbf{X}$ has to be preceded by a strategy quantifier $\langle\!\langle A\rangle\!\rangle$. Several other fragments of $\mathsf{ATL}^\star$ are also classically defined:

- ATL^+ is the restriction of $\mathsf{ATL}^\star$ where a strategy quantifier $\langle\!\langle A\rangle\!\rangle$ has to be inserted between two embedded temporal modalities $\mathbf{U}$ or $\mathbf{X}$ but boolean combination are allowed.
- EATL extends ATL by allowing the operators $\langle\!\langle A\rangle\!\rangle\,\mathbf{G}\,\mathbf{F}$ (often denoted as $\langle\!\langle A\rangle\!\rangle\,\overset{\infty}{\mathbf{F}}$) and $\langle\!\langle A\rangle\!\rangle\,\mathbf{F}\,\mathbf{G}$ (often written $\langle\!\langle A\rangle\!\rangle\,\overset{\infty}{\mathbf{G}}$). They are especially useful to express fairness properties.

For instance,

$$\begin{aligned}
\langle\!\langle A\rangle\!\rangle \Big(\mathbf{F}\,P_1 \wedge \mathbf{F}\,P_2 \wedge P_3\,\mathbf{U}\,P_4\Big) &\text{ is in } \mathsf{ATL}^+,\\
\langle\!\langle A\rangle\!\rangle\,\mathbf{F}\Big(P_1 \wedge \langle\!\langle A'\rangle\!\rangle\,\overset{\infty}{\mathbf{F}}\,P_2\Big) &\text{ is in } \mathsf{EATL},\\
\langle\!\langle A\rangle\!\rangle \Big(\mathbf{F}\,P_1 \wedge P_2\,\mathbf{U}\,(P_3 \wedge \mathbf{F}\,P_4)\Big) &\text{ is in } \mathsf{ATL}^\star.
\end{aligned}$$

3.4.1. *Model checking ATL^+.* First note that ATL^+ extends ATL and allows to express properties with more succinct formulae [Wil99, AI01] but these two logics have the same expressive power: every ATL^+ formula can be translated into an equivalent ATL formula [HRS02].

The complexity of model checking ATL^+ over ATSs has been settled Δ_3^{P}-complete in [Sch04]. But the Δ_3^{P}-hardness proof of [Sch04] is in $\mathsf{LOGSPACE}$ only w.r.t. the DAG-size of the formula. Below, we prove that model checking ATL^+ is Δ_3^{P}-complete (with the classical definition of the size of a formula) for our three kinds of game structures.

Proposition 3.14. *Model checking ATL^+ can be achieved in Δ_3^{P} on implicit CGSs.*

Proof. A Δ_3^{P} algorithm is given in [Sch04] for explicit CGSs. We extend it to handle implicit CGSs: for each subformula of the form $\langle\!\langle A\rangle\!\rangle\,\varphi$, guess (state-based) strategies for players in A. In each state, the choices of each player in A can be replaced in the transition functions. We then want to compute the set of states where the CTL^+ formula $\mathbf{A}\varphi$ holds. This can be achieved in Δ_2^{P} [CES86, LMS01], but requires to first compute the possible transitions in the remaining structure, *i.e.*, to check which of the transition formulae are satisfiable. This is done by a polynomial number of independent calls to an NP oracle, and thus does not increase the complexity of the algorithm. □

Proposition 3.15. *Model checking ATL^+ on turn-based two-player explicit CGSs is Δ_3^{P}-hard.*

Proof. This reduction is a quite straightforward extension of the one presented in [LMS01] for CTL^+. In particular, it is quite different from the previous reductions, since the boolean formulae are now encoded in the ATL^+ formula, and not in the model.

We encode an instance $\mathcal{I}$ of SNSAT_2, keeping the notations used in the proofs of Prop. 3.5 (for the SNSAT_2 problem) and 3.10 (for clause numbering). Fig. 3 depicts the turn-based two-player CGS $\mathcal{C}$ associated to $\mathcal{I}$. States s_1 to s_m are labeled by atomic proposition s,

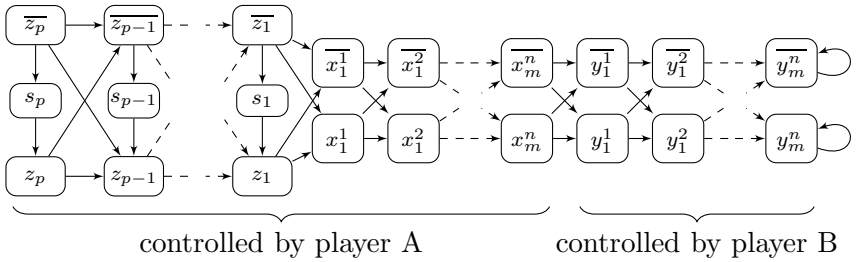

Figure 3: The CGS $\mathcal{C}$

states $\overline{z_1}$ to $\overline{z_m}$ are labeled by atomic proposition $\overline{z}$, and the other states are labeled by their name as shown on Fig. 3.

The ATL^+ formula is built recursively, with $\psi_0 = \top$ and

$$\psi_{k+1} = \langle\!\langle A \rangle\!\rangle \left[\mathbf{G}\,\neg s \wedge \mathbf{G}\,(\overline{z} \rightarrow \mathbf{EX}\,(s \wedge \mathbf{EX}\,\neg\psi_k)) \wedge \bigwedge_{w \leq p} [(\mathbf{F}\,z_w) \rightarrow \bigwedge_{j \leq n} \bigvee_{k \leq 3} \mathbf{F}\,l_w^{j,k}]\right]$$

where $l_w^{j,k} = v$ when $s_w^{j,k} = v$ and $\alpha_w^{j,k} = 1$, and $l_w^{j,k} = \overline{v}$ when $s_w^{j,k} = v$ and $\alpha_w^{j,k} = 0$. We then have:

Lemma 3.16. *For any $r \leq p$ and $t \geq r$, the following statements are equivalent:*
(a) $z_r \models \psi_t$;
(b) the strategies $f_{\mathcal{I}}$ witness the fact that $q_r \models \psi_t$;
(c) variable z_r evaluates to true in $v_{\mathcal{I}}$.

When $r = 1$, since no s- or $\overline{z}$-state is reachable from z_1, the fact that $z_1 \models \psi_t$, with $t \geq 1$, is equivalent to $z_1 \models \langle\!\langle A \rangle\!\rangle \bigwedge_j \bigvee_k \mathbf{F}\,l_1^{j,k}$. This in turn is equivalent to the fact that z_1 evaluates to true in $\mathcal{I}$.

We now turn to the inductive case. If $z_{r+1} \models \psi_{t+1}$ with $t \geq r$, consider a strategy for A s.t. all the outcomes satisfy the property, and pick one of those outcomes, say ρ. Since it cannot run into any s-state, it defines a valuation v_ρ for variables z_1 to z_{r+1} and x_1^1 to x_m^n in the obvious way. Each time the outcome runs in some $\overline{z_u}$-state, it satisfies $\mathbf{EX}\,(s \wedge \mathbf{EX}\,\psi_t)$. Each time it runs in some z_u-state, the suffix of the outcome witnesses formula ψ_{t+1} in z_u. Both cases entail, thanks to the i.h., that $v_\rho(z_u) = v_{\mathcal{I}}(z_u)$ for any $u < r + 1$. Now, the subformula $\bigwedge_w [(\mathbf{F}\,z_w) \rightarrow \bigwedge_{j \leq n} \bigvee_{k \leq 3} \mathbf{F}\,l_w^{j,k}$, when $w = r + 1$, entails that φ_{r+1} is indeed satisfied whatever the values of the y_{r+1}^j's, *i.e.*, that z_{r+1} evaluates to true in $\mathcal{I}$.

Conversely, if z_r evaluates to true, then strategy $f_{\mathcal{I}}$ clearly witnesses the fact that ψ_t holds in state z_r. □

As an immediate corollary, we end up with:

Theorem 3.17. *Model checking ATL^+ is Δ_3^{P}-complete on ATSs as well as on explicit CGSs and implicit CGSs.*

3.4.2. *Model checking EATL.* In the classical branching-time temporal logics, adding the modality $\mathbf{E}\overset{\infty}{\mathbf{F}}$ to CTL increases its expressive power (see [Eme90]), this is also true when considering alternating-time temporal logics, as we will see in Section 4.2.2.

From the theoretical-complexity point of view, there is no difference between ATL and EATL:

Theorem 3.18. *Model checking EATL is:*
- PTIME*-complete over explicit CGSs;*
- Δ_2^{P}*-complete over ATSs;*
- Δ_3^{P}*-complete over implicit CGSs.*

Proof. We extend the model-checking algorithm for ATL. This is again achieved by expressing modalities $\langle\!\langle A\rangle\!\rangle \overset{\infty}{\mathbf{F}}$ and $\langle\!\langle A\rangle\!\rangle \overset{\infty}{\mathbf{G}}$ as fixpoint formulas [dAHM01]:

$$\langle\!\langle A\rangle\!\rangle \overset{\infty}{\mathbf{F}}\, p \equiv \nu y.\mu x.\,(\langle\!\langle A\rangle\!\rangle \mathbf{X}\,(x) \vee (p \wedge \langle\!\langle A\rangle\!\rangle \mathbf{X}\,(y)))$$
$$\langle\!\langle A\rangle\!\rangle \overset{\infty}{\mathbf{G}}\, p \equiv \mu y.\nu x.\,(\langle\!\langle A\rangle\!\rangle \mathbf{X}\,(x) \wedge (p \vee \langle\!\langle A\rangle\!\rangle \mathbf{X}\,(y)))$$

Computing these fixpoints can again be achieved by a polynomial number of computations of *CPre*.

Hardness directly follows from the hardness of ATL model checking. □

3.4.3. *ATL* model-checking.* When considering ATL* model checking, the complexity is the same for explicit CGS, implicit CGS and ATS since it mainly comes from the formula to be checked:

Theorem 3.19. *Model checking ATL* is* 2EXPTIME*-complete on ATSs as well as on explicit CGSs and implicit CGSs.*

Proof. We extend the algorithm of [AHK02]. This algorithm recursively labels each location with the subformulae it satisfies. Formulas $\langle\!\langle A\rangle\!\rangle \psi$, with $\psi \in \mathsf{LTL}$, are handled by building a deterministic Rabin tree automaton $\mathcal{A}_\psi$ for ψ, and a Büchi tree automaton $\mathcal{A}_{\mathcal{C},A}$ recognizing trees corresponding to the sets of outcomes of each possible strategy of coalition A in the structure $\mathcal{C}$. We refer to [AHK02] for more details on the whole proof, and only focus on the construction of $\mathcal{A}_{\mathcal{C},A}$.

The states of $\mathcal{A}_{\mathcal{C},A}$ are the states of $\mathcal{C}$. From location ℓ, there are as many transitions as the number of possible joint moves $m = (m_{A_i})_{A_i \in A}$ of coalition A. Each transition is a set of states that should appear at the next level of the tree. Formally, given $p \in 2^{\mathsf{AP}}$,

$$\delta(\ell, p) = \{\mathsf{Next}(\ell, A, m) \mid m = (m_{A_i})_{A_i \in A} \text{ with } \forall A_i \in A.\ m_{A_i} \in \mathsf{Mov}(\ell, A_i)\}$$

when $p = \mathsf{Lab}(\ell)$, and $\delta(\ell, p) = \varnothing$ otherwise.

For explicit CGSs, this transition function is easily computed in polynomial time. For ATSs and implicit CGSs, the transition function is computed by enumerating the (exponential) set of joint moves of coalition A (computing $\mathsf{Next}(\ell, A, m)$ is polynomial once the joint move is fixed).

Computing $\mathcal{A}_{\mathcal{C},A}$ can thus be achieved in exponential time. Testing the emptiness of the product automaton then requires doubly-exponential time. The whole algorithm thus runs in 2EXPTIME. The lower bound directly follows from the lower bound for explicit CGSs. □

Let us finally mention that our results could easily be lifted to Alternating-time μ-calculus (AMC) [AHK02]: the PTIME algorithm proposed in [AHK02] for explicit CGSs, which again consists in a polynomial number of computations of *CPre*, is readily adapted to ATSs and implicit CGSs: as a result, model checking the alternation-free fragment has the same complexities as model checking ATL, and model checking the whole AMC is in EXPTIME for our three kinds of models.

4. Expressiveness

We have seen that the ability of quantifying over the possible strategies of the agents increases the complexity of model checking and makes the analysis more difficult.

We now turn to expressivity issues. We first focus on translations between our different models (explicit CGS, implicit CGS and ATS). We then consider the expressiveness of "Until" and "Always" modalities, proving that they cannot express the dual of "Until".

4.1. Comparing the expressiveness of CGSs and ATSs.

We prove in this section that CGSs and ATSs are closely related: they can model the same concurrent games. In order to make this statement formal, we use the following definition:

Definition 4.1 ([AHKV98])**.** Let $\mathcal{A}$ and $\mathcal{B}$ be two models of concurrent games (either ATSs or CGSs) over the same set Agt of agents. Let $R \subseteq \mathsf{Loc}_{\mathcal{A}} \times \mathsf{Loc}_{\mathcal{B}}$ be a (non-empty) relation between states of $\mathcal{A}$ and states of $\mathcal{B}$. That relation is an *alternating bisimulation* when, for any $(\ell, \ell') \in R$, the following conditions hold:

- $\mathsf{Lab}_{\mathcal{A}}(\ell) = \mathsf{Lab}_{\mathcal{B}}(\ell')$;
- for any coalition $A \subseteq \mathsf{Agt}$, we have

$$\forall m\colon A \to \mathsf{Mov}_{\mathcal{A}}(\ell, A).\ \exists m'\colon A \to \mathsf{Mov}_{\mathcal{B}}(\ell', A).$$
$$\forall q' \in \mathsf{Next}(\ell', A, m').\ \exists q \in \mathsf{Next}(\ell, A, m).\ (q, q') \in R.$$

- symmetrically, for any coalition $A \subseteq \mathsf{Agt}$, we have

$$\forall m'\colon A \to \mathsf{Mov}_{\mathcal{B}}(\ell', A).\ \exists m\colon A \to \mathsf{Mov}_{\mathcal{A}}(\ell, A).$$
$$\forall q \in \mathsf{Next}(\ell, A, m).\ \exists q' \in \mathsf{Next}(\ell', A, m').\ (q, q') \in R.$$

where $\mathsf{Next}(\ell, A, m)$ is the set of locations that are reachable from ℓ when each player $A_i \in A$ plays $m(A_i)$.

Two models are said to be alternating-bisimilar if there exists an alternating bisimulation involving all of their locations.

With this equivalence in mind, ATSs and CGSs (both implicit and explicit ones) have the same expressive power[4]:

Theorem 4.2. *(1) Any explicit CGS can be translated into an alternating-bisimilar implicit one in linear time;*
(2) Any implicit CGS can be translated into an alternating-bisimilar explicit one in exponential time;
(3) Any explicit CGS can be translated into an alternating-bisimilar ATS in cubic time;

[4]The translations between ATSs and explicit CGSs was already mentionned in [GJ04].

(4) Any ATS can be translated into an alternating-bisimilar explicit CGS in exponential time;
(5) Any implicit CGS can be translated into an alternating-bisimilar ATS in exponential time;
(6) Any ATS can be translated into an alternating-bisimilar implicit CGS in quadratic time;

Figure 4 summarizes those results. From our complexity results (and the assumption that the polynomial-time hierarchy does not collapse), the costs of the above translations is optimal.

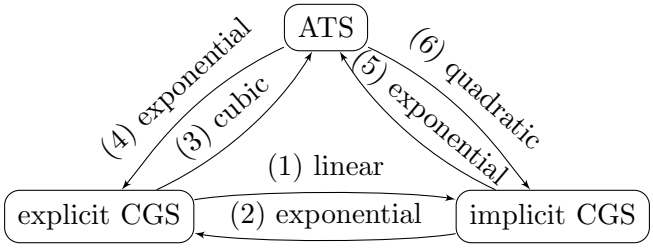

Figure 4: Costs of translations between the three models

Proof. Points 1, 2, and 4 are reasonnably easy.

For point 6, it suffices to write, for each possible next location, the conjunction (on each agent) of the disjunction of the choices that contain that next location. For instance, if we have $\mathsf{Mov}_{\mathcal{A}}(\ell_0, A_1) = \{\{\ell_1, \ell_2\}, \{\ell_1, \ell_3\}\}$ and $\mathsf{Mov}_{\mathcal{A}}(\ell_0, A_2) = \{\{\ell_2, \ell_3\}, \{\ell_1\}\}$ in the ATS $\mathcal{A}$, then each player will have two choices in the associated CGS $\mathcal{B}$, and

$$\mathsf{Edg}_{\mathcal{B}}(\ell_0) = \begin{pmatrix} (A_1 = 1 \vee A_1 = 2) \wedge (A_2 = 2), & \ell_1 \\ (A_1 = 1) \wedge (A_2 = 1), & \ell_2 \\ (A_1 = 2) \wedge (A_2 = 1), & \ell_3 \end{pmatrix}$$

Formally, let $\mathcal{A} = (\mathsf{Agt}, \mathsf{Loc}_{\mathcal{A}}, \mathsf{AP}, \mathsf{Lab}_{\mathcal{A}}, \mathsf{Mov}_{\mathcal{A}})$ be an ATS. We then define $\mathcal{B} = (\mathsf{Agt}, \mathsf{Loc}_{\mathcal{B}}, \mathsf{AP}, \mathsf{Lab}_{\mathcal{B}}, \mathsf{Mov}_{\mathcal{B}}, \mathsf{Edg}_{\mathcal{B}})$ as follows:

- $\mathsf{Loc}_{\mathcal{B}} = \mathsf{Loc}_{\mathcal{A}}$, $\mathsf{Lab}_{\mathcal{B}} = \mathsf{Lab}_{\mathcal{A}}$;
- $\mathsf{Mov}_{\mathcal{B}} \colon \ell \times A_i \to [1, |\mathsf{Mov}_{\mathcal{A}}(\ell, A_i)|]$;
- $\mathsf{Edg}_{\mathcal{B}}$ is a function mapping each location ℓ to the sequence $((\varphi_{\ell'}, \ell'))_{\ell' \in \mathsf{Loc}_{\mathcal{A}}}$ (the order is not important here, as the formulas will be mutually exclusive) with

$$\varphi_{\ell'} = \bigwedge_{A_i \in \mathsf{Agt}} \left(\bigvee_{\substack{\ell' \text{ appears in the } j\text{-th} \\ \text{set of } \mathsf{Mov}_{\mathcal{A}}(\ell, A_i)}} A_i \stackrel{?}{=} j \right)$$

Computing $\mathsf{Edg}_{\mathcal{B}}$ requires quadratic time (more precisely $O(|\mathsf{Loc}_{\mathcal{A}}| \times |\mathsf{Mov}_{\mathcal{A}}|)$). It is now easy to prove that the identity $\mathrm{Id} \subseteq \mathsf{Loc}_{\mathcal{A}} \times \mathsf{Loc}_{\mathcal{B}}$ is an alternating bisimulation, since there is a direct correspondance between the choices in both structures.

We now explain how to transform an explicit CGS into an ATS, showing point 3. Let $\mathcal{A} = (\mathsf{Agt}, \mathsf{Loc}_{\mathcal{A}}, \mathsf{AP}, \mathsf{Lab}_{\mathcal{A}}, \mathsf{Mov}_{\mathcal{A}}, \mathsf{Edg}_{\mathcal{A}})$ be an explicit CGS. We define the ATS $\mathcal{B} = (\mathsf{Agt}, \mathsf{Loc}_{\mathcal{B}}, \mathsf{AP}, \mathsf{Lab}_{\mathcal{B}}, \mathsf{Mov}_{\mathcal{B}})$ as follows (see Figure 5 for more intuition on the construction):

- $\mathsf{Loc}_{\mathcal{B}} \subseteq \mathsf{Loc}_{\mathcal{A}} \times \mathsf{Loc}_{\mathcal{A}} \times \mathbb{N}^k$, where $k = |\mathsf{Agt}|$, with $(\ell, \ell', m_{A_1}, \ldots, m_{A_k}) \in \mathsf{Loc}_{\mathcal{B}}$ iff $\ell = \mathsf{Edg}_{\mathcal{A}}(\ell', m_{A_1}, \ldots, m_{A_k})$;
- $\mathsf{Lab}_{\mathcal{B}}(\ell, \ell', m_{A_1}, \ldots, m_{A_k}) = \mathsf{Lab}_{\mathcal{A}}(\ell)$;
- From a location $q = (\ell, \ell', m_{A_1}, \ldots, m_{A_k})$, player A_j has $|\mathsf{Mov}_{\mathcal{A}}(\ell, A_j)|$ possible moves:

$$\mathsf{Mov}_{\mathcal{B}}(q, A_j) = \Big\{ \big\{ (\ell'', \ell, m'_{A_1}, \ldots, m'_{A_j} = i, \ldots, m'_{A_k}) \;\mid\; m'_{A_n} \in \mathsf{Mov}_{\mathcal{A}}(\ell, A_n) \text{ and } \ell'' = \mathsf{Edg}_{\mathcal{A}}(\ell, m_{A_1}, \ldots, m_{A_j} = i, \ldots, m_{A_k}) \big\} \;\mid\; i \in \mathsf{Mov}_{\mathcal{A}}(\ell, A_j) \Big\}$$

This ATS is built in time $O(|\mathsf{Loc}_{\mathcal{A}}|^2 \cdot |\mathsf{Edg}_{\mathcal{A}}|)$. It remains to show alternating bisimilarity between those structures. We define the relation

$$R = \{(\ell, (\ell, \ell', m_{A_1}, \ldots, m_{A_k})) \;\mid\; \ell \in \mathsf{Loc}_{\mathcal{A}}, (\ell, \ell', m_{A_1}, \ldots, m_{A_k})) \in \mathsf{Loc}_{\mathcal{B}}\}.$$

It is now only a matter of bravery to prove that R is an alternating bisimulation between $\mathcal{A}$ and $\mathcal{B}$.

Point 5 is now immediate (through explicit CGSs), but it could also be proved in a similar way as point 3. □

Let us mention that our translations are optimal (up to a polynomial): our exponential translations cannot be achieved in polynomial time because of our complexity results for ATL model-checking. Note that it does not mean that the resulting structures must have exponential size.

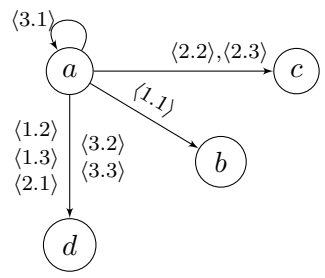

Moves from location A:

Player 1	
move 1:	$\{b_{a,1,1}, d_{a,1,2}, d_{a,1,3}\}$
move 2:	$\{c_{a,2,2}, c_{a,2,3}, d_{a,2,1}\}$
move 3:	$\{a_{a,3,1}, d_{a,3,2}, d_{a,3,3}\}$

Player 2	
move 1:	$\{a_{a,3,1}, b_{a,1,1}, d_{a,2,1}\}$
move 2:	$\{c_{a,2,2}, d_{a,1,2}, d_{a,3,2}\}$
move 3:	$\{c_{a,2,3}, d_{a,1,3}, d_{a,3,3}\}$

Figure 5: Converting an explicit CGS into an ATS

4.2. **Some remarks on the expressiveness of ATL.**

4.2.1. $\langle\!\langle A \rangle\!\rangle\, \mathbf{R}$ *cannot be expressed with* $\langle\!\langle A \rangle\!\rangle\, \mathbf{U}$ *and* $\langle\!\langle A \rangle\!\rangle\, \mathbf{G}$. In the original papers defining ATL [AHK97, AHK02], the syntax of that logic was slightly different from the one we used in this paper: following classical definitions of the syntax of CTL, it was defined as:

$$\begin{array}{rcl} \mathsf{ATL}_{\text{orig}} \ni \varphi_s, \psi_s & ::= & \top \mid p \mid \neg\varphi_s \mid \varphi_s \vee \psi_s \mid \langle\!\langle A \rangle\!\rangle\, \varphi_p \\ \varphi_p & ::= & \mathbf{X}\, \varphi_s \mid \mathbf{G}\, \varphi_s \mid \varphi_s\, \mathbf{U}\, \psi_s. \end{array}$$

Duality is a fundamental concept in modal and temporal logics: for instance, the dual of modality $\mathbf{U}$, often denoted by $\mathbf{R}$ and read *release*, is defined by $p\, \mathbf{R}\, q \stackrel{\text{def}}{\equiv} \neg((\neg p)\, \mathbf{U}\, (\neg q))$.

Dual modalities allow, for instance, to put negations inner inside the formula, which is often an important property when manipulating formulas.

In LTL, modality $\mathbf{R}$ can be expressed using only $\mathbf{U}$ and $\mathbf{G}$:

$$p\,\mathbf{R}\,q \equiv \mathbf{G}\,q \vee q\,\mathbf{U}\,(p \wedge q). \tag{4.1}$$

In the same way, it is well known that CTL can be defined using only modalities $\mathbf{EX}$, $\mathbf{EG}$ and $\mathbf{EU}$, and that we have

$$\mathbf{E}p\,\mathbf{R}\,q \equiv \mathbf{EG}\,q \vee \mathbf{E}q\,\mathbf{U}\,(p \wedge q) \qquad\qquad \mathbf{A}p\,\mathbf{R}\,q \equiv \neg\,\mathbf{E}(\neg p)\,\mathbf{U}\,(\neg q).$$

It is easily seen that, in the case of ATL, it is not the case that $\langle\!\langle A\rangle\!\rangle\, p\,\mathbf{R}\,q$ is equivalent to $\langle\!\langle A\rangle\!\rangle\,\mathbf{G}\,q \vee \langle\!\langle A\rangle\!\rangle\, q\,\mathbf{U}\,(p \wedge q)$: it could be the case that part of the outcomes satisfy $\mathbf{G}\,q$ and the other ones satisfy $q\,\mathbf{U}\,(p \wedge q)$. In fact, we prove that $\mathsf{ATL}_{\mathrm{orig}}$ is strictly less expressive than ATL:

Theorem 4.3. *There is no ATL_{orig} formula equivalent to $\Phi = \langle\!\langle A\rangle\!\rangle\,(a\,\mathbf{R}\,b)$.*

The proof of Theorem 4.3 is based on techniques similar to those used for proving expressiveness results for temporal logics like CTL or ECTL [Eme90]: we build two families of models $(s_i)_{i\in\mathbb{N}}$ and $(s'_i)_{i\in\mathbb{N}}$ s.t. (1) $s_i \not\models \Phi$, (2) $s'_i \models \Phi$ for any i, and (3) s_i and s'_i satisfy the same $\mathsf{ATL}_{\mathrm{orig}}$ formula of size less than i. Theorem 4.3 is a direct consequence of the existence of such families of models. In order to simplify the presentation, the theorem is proved for formula[5] $\Phi = \langle\!\langle A\rangle\!\rangle\,(b\,\mathbf{R}\,(a \vee b))$.

The models are described by one single inductive CGS [6] $\mathcal{C}$, involving two players. It is depicted on Fig. 6. A label $\langle\alpha,\beta\rangle$ on a transition indicates that this transition corresponds

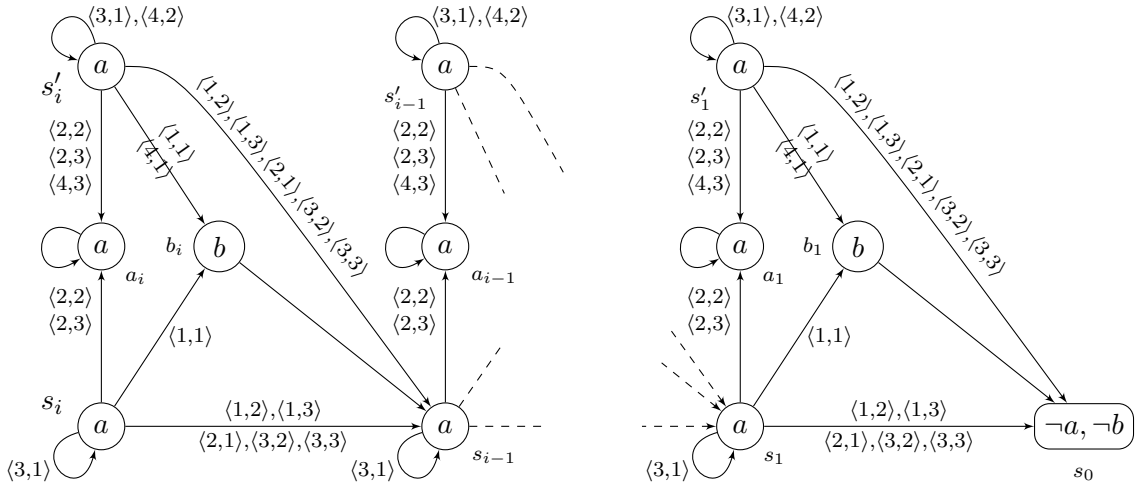

Figure 6: The CGS $\mathcal{C}$, with states s_i and s'_i on the left

to move α of player A_1 and to move β of player A_2. In that CGS, states s_i and s'_i only differ in that player A_1 has a fourth possible move in s'_i. This ensures that, from state s'_i (for any i), player A_1 has a strategy (namely, he should always play 4) for enforcing $a\,\mathbf{W}\,b$.

[5] This formula can also be written $\langle\!\langle A\rangle\!\rangle\, a\,\mathbf{W}\,b$, where $\mathbf{W}$ is the "weak until" modality.

[6] Given the translation from CGS to ATS (see Section 4.1), the result also holds for ATSs.

But this is not the case from state s_i: by induction on i, one can prove $s_i \not\models \langle\langle A_1 \rangle\rangle\, a\, \mathbf{W}\, b$. The base case is trivial. Now assume the property holds for i: from s_{i+1}, any strategy for A_1 starts with a move in $\{1,2,3\}$ and for any of these choices, player A_2 can choose a move (2, 1 and 2 resp.) that enforce the next state to be s_i where by i.h. A_1 has no strategy for $a\, \mathbf{W}\, b$.

We now prove that s_i and s'_i satisfy the same "small" formulae. First, we have the following equivalences:

Lemma 4.4. *For any $i > 0$, for any $\psi \in \mathsf{ATL}_{orig}$ with $|\psi| \leq i$:*

$$b_i \models \psi \text{ iff } b_{i+1} \models \psi \tag{4.2}$$

$$s_i \models \psi \text{ iff } s_{i+1} \models \psi \tag{4.3}$$

$$s'_i \models \psi \text{ iff } s'_{i+1} \models \psi \tag{4.4}$$

Proof. The proof proceeds by induction on i, and on the structure of the formula ψ.

Base case: $i = 1$. Since we require that $|\psi| \leq i$, ψ can only be an atomic proposition. The result is then obvious.

Induction step. We assume the result holds up to some $i - 1 \geq 1$, and prove that it then still holds for i. Let ψ s.t. $|\psi| \leq i$. We now proceed by structural induction on ψ:

- The result is again obvious for atomic propositions, as well as for boolean combinations of subformulae.
- Otherwise, the "root" combinator of ψ is a modality. If it is a CTL modality, the results are quite straightforward. Also, since there is only one transition from b_i, any $\mathsf{ATL}_{\text{orig}}$ modality can be expressed as a CTL modality in that state, and (4.2) follows.
- If $\psi = \langle\langle A_1 \rangle\rangle\, \mathbf{X}\, \psi_1$: Assume $s_i \models \psi$. Then, depending on the strategy, either b_i and s_{i-1}, or a_i and s_{i-1}, or s_i and s_{i-1}, should satisfy ψ_1. By i.h., this propagates to the next level, and the same strategy can be mimicked from s_{i+1}.

  The converse is similar (hence (4.3)), as well as the proof for (4.4).
- If $\psi = \langle\langle A_1 \rangle\rangle\, \mathbf{G}\, \psi_1$: If $s_i \models \psi$, then s_i, thus s_{i+1}, satisfy ψ_1. Playing move 3 is a strategy for player A_1 to enforce $\mathbf{G}\, \psi_1$ from s_{i+1}, since the game will either stay in s_{i+1} or go to s_i, where player A has a winning strategy.

  The converse is immediate, as player A_1 cannot avoid s_i when playing from s_{i+1}. Hence (4.3) for $\langle\langle A_1 \rangle\rangle\, \mathbf{G}$-formulae.

  If $s'_i \models \psi$, then both s'_i and s'_{i+1} satisfy ψ_1. Also, player A_1 cannot avoid the play to go in location s_{i-1}. Thus, $s_{i-1} \models \psi_1$ —and by i.h., so does s_i— and $s_i \models \psi$, as above. Now, following the same strategy in s'_{i+1} as the winning strategy of s'_i clearly enforces $\mathbf{G}\, \psi_1$. The converse is similar: it suffices to mimic, from s'_i, the strategy witnessing the fact that $s'_{i+1} \models \psi$. This proves (4.4), and concludes this case.
- If $\psi = \langle\langle A_1 \rangle\rangle\, \psi_1\, \mathbf{U}\, \psi_2$: If $s_i \models \psi$, then either ψ_2 or ψ_1 holds in s_i, thus in s_{i+1}. The former case is trivial. In the latter, player A_1 can mimic the winning strategy in s_{i+1}: the game will end up in s_i, with intermediary states satisfying ψ_1 (or ψ_2), and he can then apply the original strategy.

  The converse is obvious, since from s_{i+1}, player A_1 cannot avoid location s_i, from which he must also have a winning strategy.

  If $s'_i \models \psi$, omitting the trivial case where s'_i satisfies ψ_2, we have that $s_{i-1} \models \psi$. Also, a (state-based) strategy in s'_i witnessing ψ necessary consists in playing move 1 or 2. Thus

a_i and b_i satisfy ψ, and the same strategy (move 1 or 2, resp.) enforces $\mathbf{G}\,\psi_1$ from s_i. It is now easy to see that the same strategy is correct from s'_{i+1}. Conversely, apart from trivial cases, the strategy can again only consist in playing moves 1 or 2. In both cases, the game could end up in s_i, and then in s_{i-1}. Thus $s_{i-1} \models \psi$, and the same strategy as in s'_{i+1} can be applied in s'_i to witness ψ.

- The proofs for $\langle\langle A_2\rangle\rangle\,\mathbf{X}\,\psi_1$, $\langle\langle A_2\rangle\rangle\,\mathbf{G}\,\psi_1$, and $\langle\langle A_2\rangle\rangle\,\psi_1\,\mathbf{U}\,\psi_2$ are very similar to the previous ones. □

Lemma 4.5. *$\forall i > 0$, $\forall \psi \in \mathsf{ATL}_{orig}$ with $|\psi| \leq i$: $s_i \models \psi$ iff $s'_i \models \psi$.*

Proof. The proof proceeds by induction on i, and on the structure of the formula ψ. The case $i = 1$ is trivial, since s_1 and s'_1 carry the same atomic propositions. For the induction step, dealing with CTL modalities ($\langle\langle\varnothing\rangle\rangle$ and $\langle\langle A_1, A_2\rangle\rangle$) is also straightforward, then we just consider $\langle\langle A_1\rangle\rangle$- and $\langle\langle A_2\rangle\rangle$-modalities.

First we consider $\langle\langle A_1\rangle\rangle$-modalities. It is well-known that we can restrict to state-based strategies in this setting. If player A_1 has a strategy in s_i to enforce something, then he can follow the same strategy from s'_i. Conversely, if player A_1 has a strategy in s'_i to enforce some property, two cases may arise: either the strategy consists in playing move 1, 2 or 3, and it can be mimicked from s_i. Or the strategy consists in playing move 4 and we distinguish three cases:

- $\psi = \langle\langle A_1\rangle\rangle\,\mathbf{X}\,\psi_1$: that move 4 is a winning strategy entails that s'_i, a_i and b_i must satisfy ψ_1. Then s_i (by i.h. on the formula) and s_{i-1} (by Lemma 4.4) both satisfy ψ_1. Playing move 1 (or 3) in s_i ensures that the next state will satisfy ψ_1.
- $\psi = \langle\langle A_1\rangle\rangle\,\mathbf{G}\,\psi_1$: by playing move 4, the game could end up in s_{i-1} (*via* b_i), and in a_i and s'_i. Thus $s_{i-1} \models \psi$, and in particular ψ_1. By i.h., $s_i \models \psi_1$, and playing move 1 (or 3) in s_i, and then mimicking the original strategy (from s'_i), enforces $\mathbf{G}\,\psi_1$.
- $\psi = \langle\langle A_1\rangle\rangle\,\psi_1\,\mathbf{U}\,\psi_2$: a strategy starting with move 4 implies $s'_i \models \psi_2$ (the game could stay in s'_i for ever). Then $s_i \models \psi_2$ by i.h., and the result follows.

We now turn to $\langle\langle A_2\rangle\rangle$-modalities: clearly if $\langle\langle A_2\rangle\rangle\,\psi_1$ holds in s'_i, it also holds in s_i. Conversely, if player A_2 has a (state-based) strategy to enforce some property in s_i: If it consists in playing moves 1 or 3, then the same strategy also works in s'_i. Now if the strategy starts with move 2, then playing move 3 in s'_i has the same effect, and thus enforces the same property. □

Remark 4.6. $\mathsf{ATL}_{\text{orig}}$ and ATL have the same distinguishing power as the fragment of ATL involving only the $\langle\langle\,\cdot\,\rangle\rangle\,\mathbf{X}$ modality (see [AHKV98, proof of Th. 6]). This means that we cannot exhibit two models M and M' s.t. (1) $M \models \Phi$, (2) $M' \not\models \Phi$, and (3) M and M' satisfy the same $\mathsf{ATL}_{\text{orig}}$ formula.

Remark 4.7. In [AHK02], a restriction of CGS —the turn-based CGSs— is considered. In any location of these models (named TB-CGS hereafter), only one player has several moves (the other players have only one possible choice). Such models have the property of *determinedness*: given a set of players A, either there is a strategy for A to win some objective Φ, or there is a strategy for other players ($\mathsf{Agt}\backslash A$) to enforce $\neg\Phi$. In such systems, modality $\mathbf{R}$ can be expressed as follows: $\langle\langle A\rangle\rangle\,\varphi\,\mathbf{R}\,\psi \equiv_{\text{TB-CGS}} \neg\,\langle\langle \mathsf{Agt}\backslash A\rangle\rangle\,(\neg\varphi)\,\mathbf{U}\,(\neg\psi)$.

4.2.2. *$\langle\langle A\rangle\rangle \overset{\infty}{\mathbf{G}}$ and $\langle\langle A\rangle\rangle \overset{\infty}{\mathbf{F}}$ cannot be expressed in ATL.* It is well known that ECTL formulae of the form $\mathbf{E}\overset{\infty}{\mathbf{F}}\, P$ (and its dual $\mathbf{A}\overset{\infty}{\mathbf{G}}\, P$) cannot be expressed in CTL [Eme90]. On the other hand, the following equivalences hold:

$$\mathbf{E}\overset{\infty}{\mathbf{G}}\, P \equiv \mathbf{EF}\ \mathbf{EG}\, P \qquad\qquad \mathbf{A}\overset{\infty}{\mathbf{F}}\, P \equiv \mathbf{AG}\ \mathbf{AF}\, P.$$

The situation is again different in ATL: neither $\langle\langle A\rangle\rangle \overset{\infty}{\mathbf{F}}$ nor $\langle\langle A\rangle\rangle \overset{\infty}{\mathbf{G}}$ are expressible in ATL. Indeed, assume that $\langle\langle A\rangle\rangle \overset{\infty}{\mathbf{F}}$ could be expressed by the ATL formula Φ. This holds in particular in 1-player games (*i.e.*, Kripke structures). In the case where coalition A contains the only player, we would end up with a CTL equivalent of $\mathbf{E}\overset{\infty}{\mathbf{F}}$, which is known not to exist. A similar argument applies for $\langle\langle A\rangle\rangle \overset{\infty}{\mathbf{G}}$.

5. Conclusion

In this paper, we considered the basic questions of expressiveness and complexity of ATL. We precisely characterized the complexity of ATL, ATL$^+$, EATL and ATL* model-checking, on both ATSs and CGSs, when the number of agents is not fixed. These results complete the previously known results about these formalisms (and corrects some of them). It is interesting to see that their complexity classes (Δ_2^{P} or Δ_3^{P}) are unusual in the area of model-checking. We also showed that ATL, as originaly defined in [AHK97, AHK98, AHK02], is not as expressive as it could be expected, and we argue that the modality "Release" should be added in its definition.

References

[AHK97] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS'97)*, pages 100–109. IEEE Comp. Soc. Press, 1997.

[AHK98] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. In Willem-Paul de Roever, Hans Langmaack, and Amir Pnueli, editors, *Revised Lectures of the 1st International Symposium on Compositionality: The Significant Difference (COMPOS'97)*, volume 1536 of *LNCS*, pages 23–60. Springer, 1998.

[AHK02] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.

[AHKV98] Rajeev Alur, Thomas A. Henzinger, Orna Kupferman, and Moshe Y. Vardi. Alternating refinement relations. In Davide Sangiorgi and Robert de Simone, editors, *Proceedings of the 9th International Conference on Concurrency Theory (CONCUR'98)*, volume 1466 of *LNCS*, pages 163–178. Springer, 1998.

[AI01] Micah Adler and Neil Immerman. An $n!$ lower bound on formula size. In *Proceedings of the 16th Annual Symposium on Logic in Computer Science (LICS'01)*, pages 197–206. IEEE Comp. Soc. Press, 2001.

[CE81] Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronous skeletons using branching-time temporal logic. In Dexter Kozen, editor, *Proceedings of the 3rd Workshop on Logics of Programs (LOP'81)*, volume 131 of *LNCS*, pages 52–71. Springer, 1981.

[CES86] Edmund M. Clarke, E. Allen Emerson, and A. Prasad Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1986.

[dAHM01] Luca de Alfaro, Thomas A. Henzinger, and Rupak Majumdar. From verification to control: dynamic programs for omega-regular objectives. In *Proceedings of the 16th Annual IEEE Symposium on Logic in Computer Science (LICS'01)*, pages 279–290. IEEE Comp. Soc. Press, 2001.

[Eme90] E. Allen Emerson. Temporal and modal logic. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, chapter 16, pages 995–1072. Elsevier, 1990.

[GJ04] Valentin Goranko and Wojciech Jamroga. Comparing semantics of logics for multi-agent systems. *Synthese*, 139(2), 2004.

[GvD06] Valentin Goranko and Govert van Drimmelen. Complete axiomatization and decidability of alternating-time temporal logic. *Theoretical Computer Science*, 353(1-3):93–117, March 2006.

[HRS02] Aidan Harding, Mark Ryan, and Pierre-Yves Schobbens. Approximating ATL* in ATL. In Agostino Cortesi, editor, *Revised Papers of the 3rd International Workshop on Verification, Model Checking, and Abstract Interpretation (VMCAI'02)*, volume 2294 of *LNCS*, pages 289–301. Springer, 2002.

[JD05] Wojciech Jamroga and Jürgen Dix. Do agents make model checking explode (computationally)? In Michal Pechoucek, Paolo Petta, and László Zsolt Varga, editors, *Proceedings of the 4th International Central and Eastern European Conference on Multi-Agent Systems (CEEMAS'05)*, volume 3690 of *LNCS*. Springer, 2005.

[JD06] Wojciech Jamroga and Jürgen Dix. Model checking abilities of agents: A closer look. Technical Report IfI-06-02, Institut für Informatik, Technische Universität Clausthal, 2006.

[JD08] Wojciech Jamroga and Juergen Dix. Model checking abilities of agents: A closer look. *Theory of Computing Systems*, 42(3):366–410, 2008.

[LMO07] François Laroussinie, Nicolas Markey, and Ghassan Oreiby. On the expressiveness and complexity of ATL. In Helmut Seidl, editor, *Proceedings of the 10th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'07)*, volume 4423 of *Lecture Notes in Computer Science*, pages 243–257, Braga, Portugal, March 2007. Springer.

[LMS01] François Laroussinie, Nicolas Markey, and Philippe Schnoebelen. Model checking CTL$^+$ and FCTL is hard. In Furio Honsell and Marino Miculan, editors, *Proceedings of the 4th International Conference on Foundations of Software Science and Computation Structure (FoSSaCS'01)*, volume 2030 of *LNCS*, pages 318–331. Springer, 2001.

[Pap94] Christos Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[Pnu77] Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science (FOCS'77)*, pages 46–57. IEEE Comp. Soc. Press, 1977.

[QS82] Jean-Pierre Queille and Joseph Sifakis. Specification and verification of concurrent systems in CESAR. In Mariangiola Dezani-Ciancaglini and Ugo Montanari, editors, *Proceedings of the 5th International Symposium on Programming (SOP'82)*, volume 137 of *LNCS*, pages 337–351. Springer, 1982.

[Sch04] Pierre-Yves Schobbens. Alternating-time logic with imperfect recall. In *Proceedings of the 1st Workshop on Logic and Communication in Multi-Agent Systems (LCMAS'03)*, volume 85(2) of *ENTCS*. Elsevier, 2004.

[Wal04] Igor Walukiewicz. A landscape with games in the background. In *Proceedings of the 19th Annual Symposium on Logic in Computer Science (LICS'04)*, pages 356–366. IEEE Comp. Soc. Press, 2004.

[Wil99] Thomas Wilke. CTL$^+$ is exponentially more succinct than CTL. In C. Pandu Rangan, Venkatesh Raman, and Ramaswamy Ramanujam, editors, *Proceedings of the 19th Conferentce on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'99)*, volume 1738 of *Lecture Notes in Computer Science*, pages 110–121. Springer-Verlag, December 1999.

[WLWW06] Dirk Walther, Carsten Lutz, Frank Wolter, and Michael Wooldridge. ATL satisfiability is indeed EXPTIME-complete. *Journal of Logic and Computation*, 16(6):765–787, 2006.

This work is licensed under the Creative Commons Attribution-NoDerivs License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nd/2.0/ or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, California 94305, USA.